DIGITALEUROPE

Key Industry proposals on ErP Lot 9 draft regulation on enterprise servers and storage

Brussels, 17 March 2016

DIGITALEUROPE, the European voice of the digital technology industry, fully supports the European Commission's intentions to enhance the energy savings for servers and data storage products, as well as cost-effective improvements on materials efficiency for the circular economy. We welcome the opportunity to comment on the proposed eco-design regulation "Working document for the consultation forum on potential Ecodesign requirements for servers and data storage products". With the proposed requirements we see a high risk of missing the policy objectives.

This document suggests a number of necessary improvements and modifications to the proposed eco-design requirements in order to support the high innovation rate and continued energy and material efficiency improvements that have been achieved by the industry, and which are key to reaching the targets. For idle state power and allowances, DIGITALEUROPE will submit detailed comments separately.

EXECUTIVE SUMMARY

Impact Assessment: Measures and targets should be based on, and justified by detailed and relevant scientific assessment studies and agreed conclusions. DIGITALEUROPE would like to point to the lack of assessment and knowledge about the environmental benefits related to the proposed requirements for materials efficiency and the possible implications on intellectual property rights, product functionality and costs for the proposed requirements. Further, we note that published research on the anticipated European increase in energy consumption due to a predicted increase in data traffic and data handling shows diverging trends with regards to the growth in data centre IT equipment and the associated energy use. While some studies show increases in data centre energy use, others see stable or declining energy use based on trends in e.g. virtualization capabilities of servers and capacity optimization methods for storage products. The server and storage industry is a highly innovative, complex and fast-moving sector and we believe that the market forces are working already today to deliver the majority of improvements sought under the Lot 9 draft requirement. The continuing trend of the industry to bring to the market products which provide the capability to do substantially more work with the same or lesser material and energy investment should also be taken into account and recognized in the Commission proposal. We see no failure of market forces to address the issue properly.

Conformity assessment: Market surveillance and Harmonized standards: There is a solid common understanding between European policy makers and industry that clear, balanced and consistent conformity assessment and market surveillance procedures are important corner-stones to ensure functioning of the internal market. Therefore, we appreciate the proposed use of harmonized standards, based on appropriate standardization requests, and urge the Commission to adapt to its own timeline, e.g. for standardization under

DIGITALEUROPE
Rue de la Science, 14 - 1040 Brussels [Belgium]
T. +32 (0) 2 609 53 10 F. +32 (0) 2 431 04 89
www.digitaleurope.org | info@digitaleurope.org | @DIGITALEUROPE
Transparency register member for the Commission: 64270747023-20

M/543. Those standards should be available as soon as the requirements apply. It should be noted that harmonized standards are as important for information obligations as for e.g. technical thresholds.

Global harmonization: DIGITALEUROPE appreciates the efforts made for harmonisation of definitions, technical requirements and measurement methods, with ENERGY STAR® and other well established programs and schemes. However, we question the reference to yet unapproved and unpublished threshold limits for idle mode power consumption and capability allowances, especially in the light of ENERGY STAR® being a "top runner" program. To increase harmonization further, some definitions need to be modified and others added, and several additional product type exemptions are proposed. These changes are necessary to ensure that certain energy efficient products are not excluded from the EU market.

Materials efficiency: By far, the most significant environmental aspect related to servers and storage products is the energy consumption during use phase (~ 90 %)[1]. Given that servers and storage products operate within a B2B market the industry has well established highly efficient return, repair, reuse programmes. Today the recovery rate of storage and servers is already ~ 85 %[2]. One reason for the high recovery rate is the high content of metals, and hence high financial value. The fact that the Commission has recognized the sector as "best-in-class", the inclusion of material efficiency obligations as proposed fails to recognize the unique aspects of these products. To reiterate what we advised at the recent consultation forum, and given the lack of "significant improvement potential" and "wide disparity in the environmental performance of products available" we would like to draw your attention to potential security risks, limited usefulness at disproportionate cost, over prescriptive obligations which may hamper innovation, possible issues with intellectual property rights etc. all of which are unintended, negative consequences of the proposed regulation. DIGITALEUROPE also encourages the Commission to assess the potential energy efficiency and consumption impact of extending the life of older server and storage products as compared to promoting investment in new products which deliver more work with less energy. This balance of extending the useful life of a product versus potential energy use reductions enabled by the deployment of new product must be addressed by understanding the importance of competing political goals.

Energy Efficiency: Industry fully supports initiatives to improve energy performance but it is not clear to us how the proposed regulation will provide the savings estimated. In order to provide a meaningful assessment of server energy efficiency, DIGITALEUROPE recommends that the Lot 9 regulation remains focused on the weighted geomean active efficiency metric to collect data and establish the suitability of such a metric with regards to future for server energy efficiency requirements. DIGITALEUROPE does not believe regulating "idle mode power" as proposed is an appropriate and effective way forward. Servers are not like Personal Computing products and yet they are being compared in this way given the proposed Idle Power mode approach. Servers operate in numerous configurations and applications that impact their operational utilisation. The idle limits that are being suggested are biased toward lower performing, lower power servers. A higher performing server typically has a higher in use and idle power use, but one high performing server can do the work of many lower performing servers typically with less material, space and energy use than the smaller servers. This action also

---

1 "Ecodesign preparatory study on enterprise servers and data equipment" (47. Conclusion of task Task 7, p 386) http://bookshop.europa.eu/en/ecodesign-preparatory-study-on-enterprise-servers-and-data-equipment-pbET0415685/;pgid=GSPefJMEtXBSR0dT6jbGakZD0000pzPbSA2k;sid=DrBWSbchLItWTODDM2LI7tUECWNQmrDyCGM=?CatalogCategoryID=CXoKABst5TsAAAEjepEY4e5L Catalogue number: ET-04-15-685-EN-N

2 Working document on: Potential Ecodesign requirements for servers and data storage products, explanatory notes

demonstrates a lack of understanding of why there are high performance products offered to the market. The result of this requirement will essentially preclude the sale and deployment of many high performance servers and de facto lead to a net increase in energy consumption given that the work that is needed to be done remains, and will require more servers resulting in total combined higher power consumption.

Further, DIGITALEUROPE opposes the decision to require idle testing at a higher boundary temperature (Annex II 3.1.j) as it requires extra testing cost without providing meaningful, additional information about a typical server´s energy use or performance.

Timeline: Server and Storage products are extremely complex and involve sophisticated interactions between software and hardware. Major product transitions typically take 2-4 years minimum to accomplish. Given the expected implementation time of a possible regulation and the long design cycles of this product group, the January 1, 2019 date for the Tier 1 requirements seems too aggressive. The effective date of the Tier 1 requirements should be 2 years from OJEU publication of the implementing measure whichever is the earliest.

A. Estimated future energy consumption and possible savings - Whereas clauses (5) and (6)

The estimated annual energy consumption related to servers and storage products has been estimated and communicated in "whereas 5" of the Commission working document presented at the consultation forum. "Whereas 6" suggests the estimated energy savings potential as a direct result of eco-design measures. Whilst the figures themselves will have a degree of variation it is not entirely clear how such regulation will in fact provide the savings that are estimated and the assumptions made to deliver the savings?

Recent work conducted in the UK in support of a Climate Change Agreement for the UK commercial (colocation) data centre sector provides a good insight into the power used. It is also noted that this is auditable reported data. The audited data shows 2.15 TWh of annual power consumed by the commercial data centre sector with a similar consumption by the enterprise (in house) part of the sector. The part of the sector that supports distributed forms of IT via small data centres and server rooms is possibly accounting for 50% of the total sector consumption. The sector is estimated to use between 2 and 3% of the UK electricity consumption (which is less than 1% of total energy consumption).

DIGITALEUROPE projects that the growth of data centre consumption will occur at a slower pace or perhaps at a flat rate as opposed to the higher energy consumption growth rates projected by DG-GROW and the NGOs present at the stakeholder meeting. The DIGITALEUROPE viewpoint is based on key market trends that should be taken into account by the Commission:

1. The continued consolidation of enterprise server operations through virtualization, consolidation on true cloud platforms and real-time data centre workload management software is advancing the sector towards virtualization and more efficient operation via private and public cloud computing services.

2. Continued server performance increases in the same power envelope with each new product generation.

3. The continued propagation of cooling best practices and energy stewardship in legacy datacentres and the increased energy efficiency of new, purpose built facilities with PUE's of 1.1 to 1.3; a 50 – 70% improvement over PUE's in legacy data centres.

This viewpoint is backed up by data. A study conducted in Sweden by Telia and Ericsson[3] supports the DIGITALEUROPE viewpoint, concluding that data centre energy use in Sweden has started to decrease despite a continued exponential increase in data traffic. The November 2016 IDC Server forecast for 30 countries in Europe indicate that server unit sales will increase at an average rate of 1.6%/year, remaining largely flat for the 4 year period. Given the typical three to five year first life of a server product, the limited growth in server sales volumes indicates that the installed base will remain largely flat through 2020 when considering that there will be some virtualization and consolidation of work from multiple older servers onto a single new server. This will likely result in minimal increases in data centre energy use and consumption over that time period.

Further data from the US Data Centre Energy Usage report[4] indicates that total energy use for all data centres in the US starting from ~ 2008 was approximately 65 TWh. This estimate includes ALL servers, data storage, networking equipment, and associated infrastructure energy use in the US. Projecting data centre energy consumption through 2020, the report projected that growth would be minimal. The projected energy consumption figures within the EU would appear to be high by comparison to the US consumption. It is also very difficult to project energy consumption beyond 2020 with any degree of accuracy due to the rapid evolution of data centre IT equipment capabilities.

The latest CBRE European Data Centre market review[5] suggests that 2016 was a record breaking year for the growth of IT operations at new, higher efficiency co-location centres with demand for co-location space up over 200%. It concludes that consolidation of data centre operations into these new, more efficient data centre locations is a major contributor to managing data centre demand. Servers themselves are improving their energy efficiency with each new generation product, as will be shown in the idle power review that DIGITALEUROPE will submit to DG-GROW by March 31, 2017. The Industry continues to deliver significant improved performance with the same energy use and consumption with each new product generation over the past eight years.

In conclusion whilst we share the view that energy efficiency remains a key objective it is questionable that the impact of eco-design measures as suggested in the working document is the sole contributing factor for the estimated energy consumption savings. DIGITALEUROPE believes that it is a combination of the factors discussed above. These factors have not yet been adequately represented in the "whereas" statements in the Lot 9 draft. We would propose that factors such as virtualization, migration from in-house IT to cloud computing services and improved server performance and energy efficiency should be recognized and listed in

3 Malmodin, J., Lundén, D. (2016). The energy and carbon footprint of the ICT and E&M sector in Sweden 1990–2015 and beyond. Paper published and presented at: ICT for Sustainability (ICT4S), Amsterdam, Netherlands, 30–31 August 2016.

4 Shehabi, Arman, Sarah Smith, Dale A. Sartor, Richard E. Brown, Magnus Herrlin, Jonathan G. Koomey, Eric R. Masanet, Nathaniel Horner, Inês Lima Azevedo, and William Lintner. 2016. United States Data Center Energy Usage Report. Berkeley, CA: Lawrence Berkeley National Laboratory. LBNL-1005775. June 27. [http://eta.lbl.gov/publications/united-states-data-center-energy-usag]

5 https://f.tlcollect.com/fr2/517/35497/2016_Q4_European_Data_Centre_Marketview.pdf

the Commission document as contributing factors that are already enabling reduced energy consumption and greater efficiency.

B. Definitions and exemptions - Article 1, Article 2 and Annex I

DIGITALEUROPE has reviewed the definitions and recommends the following changes. Definition modifications are proposed to clarify technical requirements. Additional definitions of server types are proposed for exclusion, as these server types are specialty servers with limited market share, specific configurations and utilize application software which do not lend themselves to testing and assessment under the standard test method.

Definition Modifications:

Solid State Drive (SSD): this definition needs to be aligned with HDD. 'Solid State Drive' (SSD) means the primary computer storage device which reads and writes to non-volatile solid state memory devices instead of rotating magnetic platters for data storage;

Small Storage Product: The maximum number of allowable drives should be increased from 1 (in the proposed definition) to 4. Many consumer products such as USB disk drives and consumer storage systems can have up to four drives to add reliability, capacity or enable a mixed SSD/HDD system. Leaving the maximum drive count for small storage systems will bring these consumer drives under these requirements.

Low-end and High-end configuration definitions: The allowable memory levels for the two test configurations need to be modified. SPEC SERT® only supports servers which have at least one memory DIMM populated per memory channel, and have the same amount of memory populated per memory channel. Deployed servers almost all have their memory populated in this way in a datacentre. The current draft of Ecodesign DG Grow Enterprise Server Lot 9, in a number of cases, does not allow a server to be configured in this manner. DIGITALEUROPE recommends the current memory size requirements be updated to round up to the nearest multiple of, minimum supported DIMM size, multiplied by, number of memory channels. This would ensure all servers can be configured so that SPEC SERT® will function properly, and in a configuration consistent with real world use.

DIGITALEUROPE has the following concerns about the current proposal in the Draft Regulation:

1) In the case of CPUs with 8 or fewer hardware threads per socket, the Low-End configuration requirement would be 6GB per socket or less. This is in violation of SPEC SERT's minimum memory size requirement of 8GB, meaning that SERT will not function correctly.

2) SPEC's Run and Reporting rules for SPEC SERT require all memory channels to be populated with the same quantity of memory, which sometimes will not be possible under the proposed memory capacity restrictions. Additionally, in some cases where it is allowed, the user would be incentivized to add extra memory to some channels in order to use the maximum memory amount allowed, which would result in an imbalanced configuration.

3) Servers are almost never deployed in configurations which leave some memory channels unpopulated or populated in the channels with different size memory DIMMs. Servers are not designed or validated to function in this configuration.

4) Configuring a server with only some of the memory channels populated significantly reduces the system performance and causes the Operating System to handle the overhead of cores using remote memory. In one study, this reduced the performance score of the entire SPEC SERT® by more than 30%. This means, under the current proposal, even if a tester could collect a result in this unsupported configuration, the server's SPEC SERT® result would differ significantly from the actual performance of the server.

Thus, DIGITALEUROPE proposes an update to the memory size requirements which allows systems to be configured with at least one DIMM per memory channel. The following changes to the low-end and high-end configuration definitions regarding memory capacity for non-resilient and resilient servers:

High-End config "…memory capacity (in GB) equal to 3.0 times the number of total system hardware threads (rounded up to the nearest value of 8 times the number of memory channels) that represents the highest performance product model within the product family.

Low-End config: "…memory capacity (in GB) equal to 1.0 times the number of total system hardware threads (round up to the value of the GB Capacity of the minimum supported DIMM size times the number of memory channels) that represents the lowest performance product model within the product family.

A detailed analysis of the impact of the adjustments to the allowable memory capacity limits is provided in Appendix A. The allowable memory capacity multipliers are increased slightly for the high-end and low-end configurations over the Lot 9 proposal to provide manufacturers flexibility to optimize their SERT scores. Work done by the TGG SERT Analysis Working Group6 and provided during the Task Report work indicated that memory capacity beyond the SERT minimums could deliver better SERT metric scores. The higher allowable memory capacity is also appropriate, as it more closely represents the memory capacities installed in servers by data centre operators.

DIGITALEUROPE is requesting that the definitions for Solid State Drives, Small Storage Products, Low-end Performance Configuration and High-end Performance Configuration be changed as detailed above.

Exclusion for High Performance Computing (HPC) Systems, Large Servers and Storage and Network Servers Article 1 & 2:

ENERGY STAR defines HPC systems, Large Servers and Server Appliances and excludes them from the Server requirements because of the unique characteristics of these products and the fact that they are typically

6 "Analysis of the Server Efficiency Rating Tool: Implications of Server Configurations and Components on SERT™ Efficiency Results", Dietrich et.al., November 2015, Available on request from The Green Grid.

operated at a high utilization rates. The definitions and exclusions for these three product types should be included in Articles 1 and 2 of the Lot 9 regulation.

The HPC Definition reads as follows:

High Performance Computing (HPC) System means a computing system which is designed and optimized to execute highly parallel applications. HPC systems feature a large number of clustered homogeneous nodes often featuring high speed inter-processing interconnects as well as large memory capability and bandwidth. HPC systems may be purposely built, or assembled from more commonly available computer servers. HPC systems must meet ALL the following criteria: (a) marketed and sold as a Computer Server optimized for higher performance computing applications; (b) designed (or assembled) and optimized to execute highly parallel applications; (c) consists of a number of typically homogeneous computing nodes, clustered primarily to increase computational capability; and (d) includes high speed inter-processing interconnections between nodes.

Justification: HPC systems are servers utilized in large clusters targeted to maximize performance for scientific research and large scale modeling. Although some HPC clusters are based on general purpose servers, many power management features are disabled to enhance performance. Disabling power management features and the additional hardware installed significantly changes the power profile of these systems. Further, due to small market segment share in the EU, the energy consumed by HPC servers as compared to total server energy consumption is low.

As noted in the Commission's own Staff Working Document - Implementation of the Action Plan for the European High-Performance Computing strategy and the accompanying document - (COM (2016) 178 - European Cloud Initiative - Building a competitive data and knowledge economy in Europe - "HPC is the engine to power the new global digital economy, improving the scientific and industrial innovation capability and the competitiveness of industries and SMEs, allowing better services for the citizens and better decision making. HPC is one of the key contributors to the Digital Single Market (DSM) strategy next to Cloud services, Big Data and Internet of things (IoT)."

The European Commission also went on to say that "Europe is a leader in the use of HPC-powered applications: the users of HPC systems and applications in Europe include the most profitable and vibrant industrial sectors such as manufacturing, oil & gas, or pharmaceutical. HPC has excellent returns-on-investment (ROI) in Europe: for projects with financial returns, each euro invested in HPC on average returned €867 in increased revenue/income and €69 in profits."

Because of the unique configurations, high utilization and power profiles and limited market share of these products, it is appropriate to exclude HPC servers from the Lot 9 requirements.

The Large Server definition reads as follows:

'Large Server' means a resilient/scalable server which ships as a pre-integrated/pre-tested system housed in one or more full frames or racks and that includes a high connectivity I/O subsystem with a minimum of 32 dedicated I/O slots.

Several computer server manufacturers offer mainframe servers and comparable products which may be offered in 4 processor socket configurations, but which are substantially different from a non-resilient or

resilient server. While the resilient/scalable server definition underpins the definition of the large server, the system is ultimately differentiated by its increased I/O connectivity. A resilient/scalable server has a minimum of 8 dedicated I/O slots and typically cannot support more than a maximum of 16 I/O slots whereas a Large Server must have a minimum of 32 I/O slots. The power signature/characteristics of a Large Server are materially different from the power signature/characteristics of a non-resilient or resilient server with 5U or smaller enclosures, having maximum power demands of 3 kW or more. The Large Server products represent less than 0.2% of the market based on the November 2016 IDC Server forecast for 30 countries in Europe, representing a very small percentage of the energy consumed by server products. It is appropriate to exclude them from the Lot 9 requirements.

The Server Appliance definition reads as follows:

'Server Appliance' means a computer server that is bundled with a pre-installed OS and application software that is used to perform a dedicated function or set of tightly coupled functions. Server appliances deliver services through one or more networks (e.g., IP or SAN), and are typically managed through a web or command line interface. Server appliance hardware and software configurations are customized by the vendor to perform a specific task (e.g., name services, firewall services, authentication services, encryption services, and voice-over-IP (VoIP) services), and are not intended to execute user-supplied software.

Server appliances are specifically configured to enable the execution of a dedicated software system intended to perform highly specialized task(s) in real-time. They are configured to be highly available and, as such, they are maintained in a high state of readiness with minimal response times to new work commands and they are seldom, if ever, in idle mode. Server appliances also have a minimal market share with typical product volumes of several thousand products sold in a given year and they represent a very small percentage of the energy consumption of server products. It is appropriate to exclude server appliances from the scope of Lot 9.

Additionally, the Task 7 report of the preparatory study strongly recommends that the technical, economical and operational feasibility of particular eco-design measures for server appliances needs to be carefully checked. The European Commission has not yet presented any study results that check the feasibility of eco-design measures for server appliances. As server appliances are excluded from the ENERGY STAR server specification, it may very well be that these products also will need a different set of worklets and measurement methods. DIGITALEUROPE recommends to exclude server appliances from the Lot 9 requirements until further studies have been made.

DIGITAL EUROPE recommends that the Lot 9 requirements should add definitions for High Performance Server, Large Server and Server Appliance, as done by ENERGY STAR, should be included in Article 2 and the three server product types should be excluded from the Lot 9 requirements in Article 1 per the discussion above.

Definition of Server optimized for Deep Learning and Artificial Intelligence: DIGITALEUROPE has identified that servers optimized for Deep Learning and Artificial Intelligence also have unique characteristics which preclude their inclusion in the Lot 9 requirements. DIGITALEUROPE is will develop a definition detailing these servers and provide a justification for their exclusion from the Lot 9 requirements in the Idle State Power and Allowances comment document which will be submitted by 31 March 2017.

Additional Server Product Definitions and Exclusions:

Over the last 4 years, the development of software defined systems has resulted in the introduction of storage servers, network servers and hyper-converged systems. These servers are populated with large numbers of storage and/or network devices which change their fundamental power signatures and significantly increase their overall power use. They also are designed to perform specific storage and/or network functions that are not tested or evaluated by currently available server efficiency testing metrics. Because of their functional capabilities, the three "server" types are better associated with the storage and network product categories. In addition, because they are specialized systems they currently represent only a small percentage of the server market and their energy consumption is small compared to the total energy consumed by server products. DIGITALEUROPE recommends that these products be defined per the proposed definitions below and that storage servers and Hyper-converged storage systems be assigned to the on-line storage category and meet the requirements for those products, the network server be assigned to the network product category and excluded from this regulation.

Storage Server

'Storage Server' means a Storage Product which contains the same components as a computer server in addition to twelve or more storage devices. The storage server can run on more than one non-vendor specific software which is designed to support storage system connectivity, Capacity Optimization Management (COMs) Deployments and virtualized storage environments arrayed in a software defined storage network.

Network Server

'Network Server' means a network product which contains the same components as a computer server in addition to more than 11 network port with a total line rate throughput of 12 Gb/s or more, the capability to dynamically reconfigure ports and speed and support for a virtualized network environment through a software defined network.

Hyper-converged Storage System:

A Hyperconverged (HC) storage system combines properties of both compute server, data storage and network products into a single product. HC systems host client/guest instances and/or services by virtualization or containerization methods and provides the associated data storage service. Typically the systems are deployed in a cluster each sharing a portion of its direct attached storage in a scale-out manner. As such, the data storage sub-system of a HC storage appliance contains software, data storage devices, network elements, and possibly additional storage controller elements. A HC storage system is primarily accessed via network connections, and not through direct user input devices, such as a keyboard or a mouse.

Proposed definition for a Hyper-converged Storage System:

'Hyper-converged Storage System' means a Storage Product which combines the properties of compute server, data storage and network products into a single product which has one or more integrated storage tiers comprised of one or more direct attached storage products, a storage service and network methods and an integrated network fabric that enables communications between internal elements of the HC Storage System and between the HC storage system and the network.

Discussion:

Due to their extensive and specialized functions Storage, Network, and Hyper-converged Storage Systems are highly integrated devices key to future operational efficiency optimizations of the data centre. The integration of functions and the system's supporting circuitry make efficiency and power comparisons to general purpose servers, without these functions, inappropriate. The evaluation of specialized servers should be deferred until that point when a reasonable method of power and performance evaluation have been established for each of those specialized functions integrated within a general purpose server. At minimum, as long as the power use requirements for integrated versions are less than the sum of the power use for the multiple non-integrated versions, the lower power requirements of the integrated device together with its superior performance characteristics provides the more efficient option.

DIGITALEUROPE proposes that Network Servers be recognized as a network products and excluded from the Lot 9 requirements and Storage servers and Hyper-converged Storage Systems be recognized as storage products and covered by the On-line Storage Product requirements.

A. Timetable for eco-design requirements Article 3

Server and Storage products are extremely complex and involve sophisticated interactions between software and hardware. Major product transitions typically take 2-4 years minimum to accomplish. Generally compressing this type of schedule leads to quality problems and "customer line-down" situations. This is unacceptable for many customers such as banks, stock exchanges, government and airline flight scheduling enterprise customers. Given that the Lot 9 regulatory process may take until the end of 2017, an "In force" date of January 1, 2019 gives only 1 year to make design changes resulting from compliance gaps. At a minimum, the in force data should be set 2 years from publication in the OJEU to give manufacturers adequate time to address the requirements established in the regulation.

B. Power Supply Efficiency requirements. Annex II, section 1

The power supply efficiency requirements need to be modified to reflect the efficiency differences between single-output and multi-output power supplies.

1. The regulation needs to state that the efficiency requirements are for 230 Volt single output power supplies.

2. Separate efficiency requirements should be set for multi-output power supplies used in many storage products. Those power supplies carry the regulators and circuitry required to generate multiple voltage outputs for the storage product and does so more efficiently than having those regulators dispersed throughout the product. The efficiency requirements for multi-output power supplies should be set at 80plus silver for 2020 to 2022, 80plus gold for 2023 to 2025 and 80plus platinum beginning in 2026.

C. Material Efficiency requirements Annex II, section 1.2

The snapshot of slide below presented at the Consultation Forum on 17 February 2017 makes very clear that there is no environmental case for any of the proposed material efficiency requirements for servers and storage. According to Art 15.2 c of the Ecodesign directive itself "c) the product shall present significant potential for improvement in terms of its environmental impact without entailing excessive costs, …". An increase of about 1 % in recycling of servers due to the design for recycling requirement can not be considered significant. The Commission also stated that servers are best in class when it comes to recycling rates. Consequently the further condition in Art.15.2 c "ii) a wide disparity in the environmental performance of products available on the market with equivalent functionality" is also not met. If the Commission were to adhere to its own rules, all of the proposed resource efficiency requirements should be removed from the draft regulation. The products in scope in this draft regulation are B2B and their resource efficiency is dominated in the use phase. These products are unlike consumer products where many of the proposed material efficiency requirements have been developed.

Mat eff.- impact

Requirements on mat. Eff.	Design for dismantling	Securing data deletion	firmware update av.	CRM information
Can the req. can contribute to increase recycle/repair? How much?	~1% increase in recycling servers.	5~7% increase in reuse of servers and storage products.	~5% increase in reuse of servers and storage products.	~1% increase in recycling servers.
Rationale/assessment for requirements' contribution to reuse/recycle/repair	Majority of the products already comply with this. The requirement ensures the rest products will comply too, and hence increase the recycling rate very slightly.	By using data deletion software compliant with the recognised standards provides companies the confidence to allow reuse.	This requirement will remove one of the biggest barrier recyclers and second hand operators encountered currently.	more products may be disassembled for material extraction rather than completely shredded.

Source: slide presented at the Consultation Forum for servers on 17 February 2017

The fact that the Commission has recognized the sector as "best-in-class", the inclusion of material efficiency obligations as proposed fails to recognize the unique aspects of these products. Today the recovery rate of storage and servers is already ~ 85 %[7]. The server and storage products are by design modular products, configured to simplify the maintenance and replacement of individual components which in turn enables easy access to these components for recycling. To reiterate our remarks at the recent Consultation Forum, and given

7 Working document on: Potential Ecodesign requirements for servers and data storage products, explanatory notes

the lack of "significant improvement potential" and "wide disparity in the environmental performance of products available" we advise to abstain from a dismantling requirement.

D. Proposed ban on gluing and welding Annex II, section 1.2.1

The proposed draft regulation would prohibit welding and gluing of HDD and SSD, memory, processor (CPUs), motherboard, chassis, expansion cards/graphic cards, and power supply components. The industry has strong concerns about the approach to such a requirement:

- The European Commission has not presented evidence that welding and gluing inherently inhibit the removability of components, and that other fastening techniques are always preferable to achieve the objective of the requirement. The requirement further lacks credibility when the technical feasibility, benefits, costs and impact on product safety, reliability and usability have not been carefully studied. The complete absence of a thorough assessment implies the requirement has no creditable basis for inclusion. It also takes no account of the complexities of such a requirement for specific applications and assumes it can be adopted universally without exception. This was recently confirmed at the stakeholder meeting on the review of another vertical regulation for computers on 16 January 2017, where the same requirements for different components were presented by the JRC. During the meeting, the JRC stated that an assessment of technical feasibility, environmental benefits and costs of the requirements was still work in progress. We call on the European Commission to respect the process set forth in the Ecodesign Directive Art.15.4b, which calls for requirements to be backed up by data with prior assessment of the "impact on the environment, consumers and manufacturers (…) innovation, market access and costs and benefits".

- The draft wording prescribes a design trait and does not allow manufacturers to adopt innovative and efficient solutions to achieve the broad objective of the requirement, in this case removability of components for selective treatment as outlined in Directive 2012/19/EU Annex VII and referenced in Recital 21 of the draft regulation. Prohibiting the use of certain fastening techniques is too prescriptive and hampers future innovation and competitiveness of the IT industry. The product design is an industry responsibility which the regulation should respect. The focus should therefore be on the outcome to achieve the regulatory requirement. If the European Commission wants to enable removability of key components for selective treatment, the requirement should cover all joining techniques precluding removal of the components without explicitly prohibiting certain fastening techniques. The CENELEC standardization action to address material efficiency will aim to address the broad objective of removability, and as such will not be too prescriptive in the design envisaged in respect of the differing design techniques that compliment easy removal of key components.

DIGITALEUROPE recommends that the requirement prohibiting welding or firm gluing as joining or sealing techniques for the listed components used in server and storage products to be removed from the Lot 9 requirements.

E. Data deletion software for storage equipment Annex II, sections 1.2.2

Summary: DIGITAL EUROPE considers that making available built-in data deletion software inappropriate for several reasons and request this requirement to be removed. We do not perceive there being a lack of incentive for reuse and recycling, as the reuse rate of storage and servers is already ~85%, and the recycling value of products which are mainly made of metals is also very high. In addition, we have serious security concerns should such software be built-in for every product sold into the EU. This would jeopardize cyber security integrity and the relationship between manufacturers and their clients.

Management of data deletion is best done at dedicated processing locations to ensure appropriate data deletion. This is the system in place today and it is working adequately and we do not see a market failure. There is therefore no need to require installation of data deletion software on equipment. A specific data deletion methodology should not be mandated in a product energy efficiency regulation for the reasons outlined below.

Security Risk

Placing data deletion software on all storage products opens the opportunity to security breaches (hacking) which could potentially delete data remotely. This type of risk should not be encouraged by mandatory regulatory requirements.

Customer specific

Data deletion is customer specific, and there are various methods other than data deletion software for the non-destructive scrubbing or protection of storage drive data. Manufacturers, and in particular end-users, should be free to choose their approach the security of their data, whether during the use or at the end of its (first) life. In specific high security circumstances, such as the financial sector, customers will insist on the destruction of the drive to ensure data destruction.

Manufacturer's liability

If a data deletion tool does not meet customer's requirement, or fails to delete the data, there will be liability issues for the manufacturer vis-à-vis our customers.

Software license

Any software provided in commercial activities will be subject to license agreements. As with firmware there are intellectual property rights subject to contractual agreements (see under 'firmware' for further details). Placing the software on each storage product incurs a licensing cost for each system or each drive shipped in a system

DIGITALEUROPE recommends that the requirement for server and storage products be required to ship with data deletion software installed be removed from the Lot 9 requirements.

F. Information Requirements Annex II, section 3

1. Annex II Section 1, 1.1.1, 1.1.2, 1.1.3, 3.2, 3.3, & 3.4: The term "data storage product(s)" should be replaced with "Online data storage product(s)". This is consistent with the scope for Storage products present in Chart 17 of the chart deck "Brussels – 17/02/2017 Ecodesign Consultation Forum on DB GROW Lot 9 (servers and data storage products) which identifies the scope as Online 2, 3 and 4 data storage products. Our understanding of the intent of the Lot 9 regulations is to address the Emerald Taxonomy On-line 2 to On-line 4 storage product categories which represent the bulk of the storage product energy use in the data centre. When the on-line data storage product definition is combined with the large and small data storage products definitions, the remaining products are those in the On-line 2 to On-line 4 groups. If the broader definition of storage product is used in Annex II, that will draw in the following product categories:

    a. Near on-line storage: These products have minimal unit sales and minimal energy use in the market and data centre respectively.

    b. Removable Media Library: These products have a very low energy use, typically around 100 Watts as system power demand consists of a tape extractor and drive. These systems represent a very small percentage of the data centre energy consumption.

    c. Virtual Media Library: These products have minimal, if any, unit sales and no discernable energy demand in the data centre.

   Because of their low market share and/or low energy use in the data centre, these three storage product categories should be excluded from the Lot 9 requirements. The term "Online data storage product" should be used in Annex II to designate those products covered by the requirements.

2. Annex II Section 1.2.2: The term "data storage equipment" should be replaced with "data storage device", the definition provided in Article 2 for SSD and HDD devices.

3. Annex II Section 2.1: DIGITALEUROPE is developing comments and recommendations regarding the Lot 9 Draft's proposal for Idle State Power and Idle Power Allowances for servers. In an email dated 22 February 2017, the Commission granted DIGITALEUROPE an extension for the submittal on Idle State Power and Idle Power Allowances to 31 March 2017. DIGITALEUROPE will submit its comments on or before 31 March 2017.

4. Section 3.1 and 3.2: Information Provided by Manufacturers: The data required in section 3.1 and 3.2 should:

    a. The data in 3.1 (e) to (h), (j), and (k) and 3.2 (e) and (f) should only be required to be reported on the free access website of manufacturers, their authorized representatives and importers. Including the data in the instruction manuals for installers and end users has no purpose, as the data is specific to the measured configurations.

    b. The measured power use and efficiency score and the component detail should be provided for the low-end and high-end server configurations (section 3.1).

c. Section 3.1.h: We recommend that the reported maximum power for the low-end and high-end configurations be the highest maximum power recorded for the individual workloads or worklets tested by the activity efficiency test metric. The definition of "maximum power" should be added to Annex I.

d. Section 3.1.k: We recommend that the weighted geomean value for the CPU, memory and storage workloads and the overall weighted geomean server efficiency value as measured by the required active efficiency test metric be reported as server efficiency. We recommend that the interval performance and power data should not be directly reported. If the commission desires to collect all of the performance and power data measured under the specified active efficiency test metric, then a requirement should be made to make available the full test output file for all tested products and product families on a website, ideally one managed by the Commission so that the data is easily accessible.

e. Section 3.1.n: Listing all possible configurations beyond the product model is very burdensome and does not add any value to the end-user perspective. For a product model, millions of configurations are possible. The requirement should be to list the applicable model number(s) for the product family.

G. The ASHRAE table Annex II, section 3.1 (i)

Table 6 in Annex 2 needs to be modified to show the recommended and allowable temperature range for each ASHRAE environmental classes and the technical notes that explain the intent of the two ranges. The modified table and notes that should be inserted into the Lot 9 requirements are provided below.

	Dry Bulb Temperature (°C)		Humidity Range, Non Condensing		Maximum Dew Point (°C)	Maximum Rate of Change (°C/hr) (f)d
Operating condition class	**Allowable**	**Recommended**	**Recommended**	**Allowable**		
A1	15 to 32	18 to 27	-9.0oC DP to 60% RH and 15oC DP	-12oC DP to 8% RH to 80% RH	17	5/20
A2	10 to 35	18 to 27	Same as A1	Same as A1	21	5/20
A3	5 to 40	18 to 27	Same as A1	-12°C DP to 8% RH to 85% RH	24	5/20
A4	5 to 45	18 to 27	Same as A1	-12°C DP to 8% RH to 90% RH	24	5/20

Recommended Range: The temperature and humidity ranges "under which IT equipment would operate the most reliably while still achieving reasonably energy-efficient data centre operation.8

Allowable Range: The temperature and humidity ranges under which "it is acceptable to operate outside the recommended envelope for short periods of time without affecting the overall reliability and operation of the IT equipment.

H. Idle Testing at Elevated Temperatures Annex II, 3.1 (j)

DIGITALEUROPE strongly opposes any decision to require idle testing at a higher boundary temperature (Annex II 3.1.j) as it requires extra testing cost without providing meaningful, additional information about a server energy use or performance. Performing the active efficiency and idle tests at higher temperatures is unreasonably expensive as it requires the testing to occur in highly specialized environmental chambers. In addition, servers will not be operating at the higher allowable ASHRAE temperature range for an appreciable period of time. As the industry has reiterated during the various task and study periods, the ASHRAE temperature classifications involve a recommended temperature range of 18oC to 27oC, which is where the server will typically operate, and an allowable temperature range (up to 35oC for the A2 classification) where a server can operate for limited durations of time without a significant loss in reliability. Measuring idle power at the allowable range offers no meaningful data for the following reasons:

a) The intent of the higher allowable operating range is to enable data centre operators to extend the time available for free cooling. Using the most extreme condition of direct air free cooling, a server would be operating in the allowable range during the hottest periods of the day – which would likely occur from 2 pm to 6 pm. In turn, this is likely to be the period of highest server use during the day with the server experiencing little, if any, idle time. Idle periods are most likely to occur in the late evening and early morning when ambient temperatures are in the recommended temperature range.
b) Server idle power in the allowable range will increase because the cooling fans will be running at higher speeds to ensure the server is adequately cooled. In most products, a point will be reached where the extra power required to run the fans will eclipse the reduction in energy required to cool the data centre. There is little, if any, energy reduction benefit to be gained by running in the allowable temperature range for an extended period of time. As discussed in (1) above, that will not be the operating profile even in a data centre with direct air free cooling.
c) Operating at elevated temperatures for the time required to complete an active efficiency test (2 hours) or an idle test (30 minutes) is not a sufficient length of time to validate the ability of a server to stabilize its operation at that temperature.
d) Server idle power, excluding the fan power, will not vary appreciably over the recommended and allowed temperature ranges. The difference in server power use at higher temperatures will be a direct result of the higher fan speeds required to keep server operating temperatures within the

8 "ASHRAE TC 9.9 2011 Thermal Guidelines for Data Processing Environments – Expanded Data Centre Classes and Usage Guidance, published 2011, page 10.

allowable tolerances of the various system components. If it is desired to know the additional idle power driven by the fan speed, manufacturers can provide a curve of fan power to server operating temperature. Idle power at elevated temperatures can be calculated by adding the idle power at the test temperature and the additional fan power required between the test temperature and the operating temperature of interest to find the idle power at higher operating temperatures. This reasonably accurate and much more practical way to provide power use values at higher operating temperatures.

DIGITALEUROPE recommends that manufacturers be required to provide the fan power to temperature curves for the server product to enable data centre operators to calculate idle power at higher temperatures. There is no reason to require expensive, specialized testing. The requirement to test servers at the specified ASHRAE classification allowable range maximum temperature should be removed from the Lot 9 requirements.

I. Requirement for the reporting of cobalt, neodymium, and palladium Annex II, section 3.3

Summary: DIGITALEUROPE opposes the draft requirement on critical raw materials disclosure as described in Annex II 3.3, as we consider it will have a limited usefulness at a disproportionately high cost. There currently is no clear demand from recyclers to have this level of detail available for recycling purposes, in particular at a product level, as recyclers know already where the valuable materials reside in the products. The presence of the listed critical metals is dispersed in very small quantities across many components and will not be recovered as the cost of recovering the low quantities is economically not viable. In addition, it is not clear how these requirements would serve authorized 3rd parties such as upgrade/maintenance or reuse/repair organizations.

Details:

Recycler's knowledge

To our knowledge, recycler themselves determine which valuable materials they will seek to recover from waste, including electronic waste. In doing so, they will occasionally carry out their own analytical test for new types of equipment to determine the quantities of specific material/metals in order to determine the economics of recovery. While it may be possible (at a high cost) to sum up all the Palladium, Neodymium and Cobalt in a product to one decimal place, it is unlikely to be of much value regarding harvesting these metals from the product. The presence of recyclable quantities of these metals in HDD disks and magnets is known and does not require a general reporting of product content from these three metals. The presence of cobalt might be of relevance in lithium ion batteries, yet the recycling of these is covered via the Batteries Directive. We also want to emphasize that that the recovery of cobalt from batteries happens without the need of knowing the exact quantities of cobalt in the different type of lithium ion chemistries.

Low quantities of Cobalt, Palladium and Neodymium

With the exception of the HDD platters that may use cobalt as a magnetic coating, the other parts using of cobalt make up less than 1% of the components in the unit and the dispersed, alloyed nature of the parts, makes recycling impractical and not cost effective. Likewise Palladium is general found as a metallization electrode in ceramic capacitors and other parts. Extraction and recycling of this metal is

unlikely to be cost effective or practical because the Palladium is fused to the ceramic materials and typically located internally in the component. Neodymium is typically used in headphone and hard drive magnets with occasional use to attach plastic parts such as Bezels. Besides these three uses, other minor use is as a dielectric additive in small ceramic capacitors or inductors and is likewise impractical to harvest at recycle. The use of Neodymium magnets is limited to hard drive magnets in most storage devices.

Material disclosure cost:
Tracking material or chemical presence in hardware products requires significant supply chain and IT resources to gather this information at a product level. The IT sector has extensive experience with substance restriction tracking for EU RoHS purposes, as it help us to determine the compliance status of a given product. It should be noted that substance restriction tracking, which means ensuring the absence, does not require the same level of resources of presence tracking which needs to determine what, where and how much. While we can indeed leverage our existing processes, they will add cost to the existing infrastructure, yet it will not deliver any value for recyclers, and certainly not for authorized repair centres or maintenance organizations.

DIGITALEUROPE recommends that the requirements to quantify the presence of Cobalt, Palladium and Neodymium in components within the server and storage products be removed from the Lot 9 requirements.

J. Firmware
Annex II, section 1.2.3 and section 3.3(c)

Summary: DIGITALEUROPE opposes the inclusion of this requirement and request it is removed from any proposed regulation. The requirement in itself does not contribute to the material efficiency objectives, fails to recognize the high level of reuse, repair and refurbishment practices already occurring in the market, and undermines the existing firmware and refurbishment business of OEMs and their certified partners. This proposal would inflict direct commercial harm to OEMs and their certified partners, who invest significantly into the development of firmware and training, skills and service and recovery operations that ensures the firmware operates properly when installed. DIGITALEUROPE considers that there is no market failure on product reuse, and this specific proposal interferes with intellectual property rights of manufacturers and anti-trust legislation. As proposed, this requirement could jeopardize the EU's cybersecurity and data privacy policy objectives.

Details:
Intellectual property
Firmware for servers and storage products is intellectual property of the manufacturer. This software product, and associated updates, constitute a product update to hardware product. Proprietary Firmware availability is subject to contractual conditions (license) in application of the principle of freedom of contract and/or separate support offerings to customers by the manufacturer and their certified vendors. In the express agreement between the manufacturer and the customer, the warranty can stipulate the lifetime of the availability of firmware for a specific new product. The same applies to updates: availability is subject to the legal terms of use, aligned to purchasing customers and not to (uncertified) 3rd parties. Manufacturers make significant investments in the development of products and services, and the protection of intellectual property is a legitimate and important

aspect of sustaining the health of the competitive and innovative technology industry. Manufacturers oppose any mandatory licensing of their intellectual property or licensing without appropriate contractual terms unique to the nature of the licensed property and relationship of the licensor and licensee.

Data centre security

A data centre's security and privacy are fundamental goals in the design of our membership's hardware, software and services. Servers and storage products in data centres are constantly at risk from hackers and any weakening of those standards such as sharing sensitive firmware will increase risks to data security. Manufacturers also have strong concerns about independent service providers who may take risks or cut corners leaving themselves or consumers in danger if they perform service without the proper training or safety standards.

Operating system implications are not considered

At a higher level, the proposed requirement for firmware does not consider the operating system (OS) implications. For some firmware updates, OS are required to facilitate firmware tests on the functionality and compatibility of different components. In addition, OS vendors will have license agreements with the manufacturers but also separate license agreements with end-users. The requirements as currently written for firmware have not considered the above issues and it is not clear where the requirements differentiate between new products and second hand products. The repurposing of server and storage equipment is a current practice by manufacturers and our certified vendors which is bound by legal terms of use and is sometimes OS dependent.

Firmware is often customer specific

In addition, different data centre operators/customers will often choose to use a specific version/release of the firmware so they have consistency across their operation. The version/release will differ from client to client and data centre to data centre dependent on how a particular type of server is configured and used.

No market failure for product reuse

It should also be noted that this requirement is based on the need to support the informal channels. Within the electronics sector, servers and storage are high value products. This value is maintained for a long time and the total recovery rate is very high (85%), as reflected in the explanatory note accompanying the draft regulation. As such, the industry fails to see a market failure which needs to be addressed. The aims of the eco-design implementing measures are improving the environmental performance of products, without entailing excessive costs. This specific requirement in itself does not intrinsically contribute to the envisioned energy or material savings, yet potentially directly affects business models resulting in misuse of equipment and illegal copying of firmware.

DIGITALEUROPE recommends that the requirement to make the most recent firmware version must be made publicly available be removed from the Lot 9 requirements. Market mechanisms exist for all resellers and refurbishers to gain access to the firmware while protecting the intellectual property rights of the manufacturers.

K. Measurements and calculations

Annex III

a) Annex III Section 2: Modify language regarding the "as-shipped" configuration to read "Servers shall be tested in two configurations which conform to the definitions for high-end and low-end configurations in Annex II. System settings in the firmware and the operating system chosen for test shall be set in the settings that will be used for that operating system in a shipped product. The purchaser of the server selects if a server ships with an operating system, a hypervisor or no operating system (where the purchaser installs their standard operating system/hypervisor template during the installation process).

Server products can be shipped in a multitude of product configurations with any one of several operating systems and/or hypervisors or no operating system at all. Product testing has to be done on the two defined configurations otherwise the testing system becomes unworkable as every configuration would have to be tested before shipping. Any verification testing for the idle value outside of the low-end and high-end configurations will need to be conducted on a chosen product with the test operating system and settings. The measured idle value will have to be compared to the calculated idle value for the selected product based on the components used in the selected configuration.

b) No adders are listed for Auxiliary Processing Accelerators (APAs) Annex II Section 2.1 Table 5:

Removable Graphics and Expansion Cards: The draft regulation does not include idle adders for compute functions currently provided by Graphics Card and Expansion Cards. These cards/functions can include Graphical Processing Units (GPUs), Field Programmable Gate Arrays (FPGAs), and other compute intensive or specialty devices or solutions. These functional capabilities are prevalent in systems designed for high performance computing and compute intensive workloads. Currently, these cards primarily reside in a PCI slot of a server. In this configuration, the impact of APA devices offered today is easily excluded from idle and active efficiency testing by excluding the APA cards from the server product being tested.

DIGITALEUROPE proposes that a sentence be added in Annex III, section 2 after the first sentence discussing the "as-shipped" requirement which states "Graphics Cards and Expansion Cards shall not be included in any test system and shall be removed before any testing occurs."

Direct Attach APAs: APAs can also be integrated directly into the server components such as the processor motherboard. In these cases, the components add directly to the idle and operating power of the server. Currently, direct attach APAs are used on a very limited number of specialty products which are, for the most part, not able to be tested with SERT because of the application software required to integrate the GPU and the primary processor. Integrating these devices into a server systems achieves power reductions and performance improvements as compared to graphics and expansion cards, enabling more workload to be delivered per unit of energy consumed. These direct attach solutions require a significant amount of engineering and cost to implement, therefore DIGITALEUROPE is not concerned that exemption would be utilized as a loophole for shipping inefficient systems

DIGITALEUROPE proposes that a definition for Auxiliary Processing Accelerators be included in Annex II and that servers with direct attach APAs be excluded from the regulation.

L. Verification procedure for market surveillance Annex IV

Section 2 of this annex needs to be revised in one of two ways. It either needs to specify that market surveillance be performed on a product matching one of the two configurations tested to provide the data specified in Section 3.1, in which case the measured data can be checked against the reported data, or it needs to specify that the chosen configuration is tested and the idle power compared to the calculated idle power limit using the idle adders corresponding to the components in the chosen configuration. It should also be emphasized that almost all servers are built to order; servers are typically not bought "off-the-shelf".

DIGITALEUROPE recommends that one of the two tested configurations be obtained for market surveillance and tested values compared to the measured, reported values per 3.1.

--

For more information please contact:
Sylvie Feindt, DIGITALEUROPE's Sustainability Policy Director
+32 2 609 5319 or Sylvie.Feindt@digitaleurope.org

Appendix A: DIGITALEUROPE recommended updates to Family Memory Sizes for Ecodesign DG Grow Enterprise Servers

March 2017

Executive Summary

SPEC SERT® only supports servers which have at least one memory DIMM populated per memory channel, and have the same amount of memory populated in all memory channels. Deployed servers almost all have their memory populated in this way in a datacenter. The memory limitations established in the low-end and high-end configurations in the draft of Ecodesign DG Grow Enterprise Server Lot 9 do not, in a number of cases, allow a server to be configured in this manner. TGG recommends the current memory size requirements are updated to include larger multipliers against the SERT test minimum requirements and the ability to round up the allowable memory capacity to the nearest multiple of the minimum supported DIMM size times the number of system memory channels. This would ensure all servers can be configured so that SPEC SERT® will function properly and represent configurations consistent with real world use.

The need to have each memory channel populated with an equal number and capacity of DIMMs necessitates that the memory requirements for the low-end and high-end configurations be modified as follows:

Low-End: Memory size (GB) = Total System Hardware Threads * 1.0, rounded up to the nearest product of the Minimum Supported DIMM Size * Number of System Memory Channels.

Excel formula: Ceiling.Math(Hardware Threads * 1.0, Minimum Supported DIMM Size * Number of System Memory Channels)

High-End: System Hardware Threads * 3.0, rounded up to the nearest product of the Minimum Supported DIMM Size * Number of System Memory Channels)

Excel formula: Ceiling.Math(Hardware Threads * 3.0, Minimum Supported DIMM Size * Number of System Memory Channels)

The following changes to the low-end and high-end configuration definitions regarding memory capacity for non-resilient and resilient servers. The total system hardware threads are calculated by multiplying the number of processor sockets times the number of cores per socket times the number of threads per core.

High-End config "…memory capacity (in GB) equal to 3.0 times the number of total system hardware threads (rounded up to the nearest value of 8 times the number of memory channels) that represents the highest performance product model within the product family.

Low-End config: "…memory capacity (in GB) equal to 1 times the number of total system hardware threads (rounded up to the value of the GB Capacity of the minimum supported DIMM size times the number of memory channels) that represents the lowest performance product model within the product family.

Details

The current draft of EU Lot9 for Enterprise Servers specifies the family memory sizes as:

High-End config: memory capacity (in GB) equal to 1.0 to 2.0 times the product of the number of CPUs, cores and hardware threads that represents the highest performance product model within the product family.

Low-End config: memory capacity (in GB) equal to 0.5 to 0.75 times the product of the number of CPUs, cores and hardware threads that represents the lowest performance product model within the product family.

DIGITALEUROPE has the following concerns about this proposal:

1) In the case of systems with 8 or fewer hardware threads, the Low-End config requirement would be 6GB or less. This is in violation of SPEC SERT's minimum memory size requirement of 8GB, meaning that SERT will not function correctly.

2) SPEC's Run and Reporting rules for SPEC SERT require all system memory channels to be populated, which sometimes is not possible under the proposed memory capacity restrictions. Additionally, in some cases where it is allowed, the user would be incented to add extra memory to some channels in order to use the maximum memory amount allowed, which would result in an imbalanced configuration.

3) Server are almost never deployed in configurations which leave some memory channels unpopulated or populate the channels with different size memory DIMMs. Servers are not designed or as fully validated to function in this configuration.

As an example, the Intel Xeon® 2623v4 is a current generation server CPU which has 4 cores, 2 hardware threads per core, and 4 memory channels. This is a typical Low-End config CPU, and under the current proposal would only be allowed to populate up to 6GB of memory per socket. Since 4GB DIMMs are the smallest currently available, this would require a user to populate only one of the four memory channels with 1x 4GB DIMM. In actual use, this server would almost always be configured with at least 4x4GB DIMMs.

4) Configuring a server with only some of the memory channels populated significantly reduces the system performance and causes the Operating System to handle the complexity of cores using remote memory. In one study, this reduced the performance score of the entire SPEC SERT® by more than

30%. This means, under the current proposal, even if a tester could collect a result in this unsupported configuration, the server's SPEC SERT® result would differ significantly from the actual performance of the server.

Example Calculations:

The tables below show the High-End and Low-End allowed memory sizes in three scenarios. The data in the table is the maximum allowed memory amount and the horizontal (x-axis) has the number of CPU cores per socket. The calculations assume that the CPU supports two hardware threads per core.

The top table shows the maximum allowed memory amounts under the current Lot 9 proposal. The middle table shows the maximum allowed memory under the TGG proposal for a server which has 2 system memory channels and supports a minimum of 4GB memory DIMMs. The bottom table is the same as the middle table, but this time is for a system which has 4 system memory channels and supports a minimum of 8GB memory DIMMs.

As you can see, in the middle and bottom tables, the allowed memory is always equal to or more than the smallest amount required to populate every memory channel in the system with the minimum size memory DIMM.

Current Lot9 Max Memory Size Proposal:

	Memory per Socket (GB)														
Cores per socket	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32
Lot 9 Low	6	9	12	15	18	21	24	27	30	33	36	39	42	45	48
Lot 9 High	16	24	32	40	48	56	64	72	80	88	96	104	112	120	128
	* Assumes 2 hardware threads per core														

Examples under the proposal:

*If a system has **2** memory channels and the smallest supported DIMM is **4**GB, then:*

	Memory per Socket (GB) with Balanced **4GB** DIMMs and **2** Memory Channels														
Cores per socket	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32
Lot 9 Low	8	16	16	16	24	24	24	32	32	40	40	40	48	48	48
Lot 9 High	16	24	32	40	48	56	64	72	80	88	96	104	112	120	128
	* Assumes 2 hardware threads per core														

*If a system has **4** memory channels and the smallest supported DIMM is **8**GB, then:*

	Memory per Socket (GB) with Balanced **8GB** DIMMs and **4** Memory Channels														
Cores per socket	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32
Lot 9 Low	32	32	32	32	32	32	32	32	32	64	64	64	64	64	64
Lot 9 High	32	32	32	64	64	64	64	96	96	96	96	128	128	128	128
	* Assumes 2 hardware threads per core														

ABOUT DIGITALEUROPE

DIGITALEUROPE represents the digital technology industry in Europe. Our members include some of the world's largest IT, telecoms and consumer electronics companies and national associations from every part of Europe. DIGITALEUROPE wants European businesses and citizens to benefit fully from digital technologies and for Europe to grow, attract and sustain the world's best digital technology companies.

DIGITALEUROPE ensures industry participation in the development and implementation of EU policies. DIGITALEUROPE's members include 61 corporate members and 37 national trade associations from across Europe. Our website provides further information on our recent news and activities: http://www.digitaleurope.org

DIGITALEUROPE MEMBERSHIP

Corporate Members

Airbus, Amazon Web Services, AMD, Apple, BlackBerry, Bose, Brother, CA Technologies, Canon, Cisco, Dell, Dropbox, Epson, Ericsson, Fujitsu, Google, Hewlett Packard Enterprise, Hitachi, HP Inc., Huawei, IBM, Intel, iQor, JVC Kenwood Group, Konica Minolta, Kyocera, Lenovo, Lexmark, LG Electronics, Loewe, Microsoft, Mitsubishi Electric Europe, Motorola Solutions, NEC, Nokia, Nvidia Ltd., Océ, Oki, Oracle, Panasonic Europe, Philips, Pioneer, Qualcomm, Ricoh Europe PLC, Samsung, SAP, SAS, Schneider Electric, Sharp Electronics, Siemens, Sony, Swatch Group, Technicolor, Texas Instruments, Toshiba, TP Vision, VMware, Western Digital, Xerox, Zebra Technologies.

National Trade Associations

Austria: IOÖ
Belarus: INFOPARK
Belgium: AGORIA
Bulgaria: BAIT
Cyprus: CITEA
Denmark: DI Digital, IT-BRANCHEN
Estonia: ITL
Finland: TIF
France: AFNUM, Force Numérique, Tech in France
Germany: BITKOM, ZVEI
Greece: SEPE
Hungary: IVSZ
Ireland: TECHNOLOGY IRELAND
Italy: ANITEC
Lithuania: INFOBALT
Netherlands: Nederland ICT, FIAR
Poland: KIGEIT, PIIT, ZIPSEE
Portugal: AGEFE
Romania: ANIS, APDETIC
Slovakia: ITAS
Slovenia: GZS
Spain: AMETIC
Sweden: Foreningen Teknikföretagen i Sverige, IT&Telekomföretagen
Switzerland: SWICO
Turkey: Digital Turkey Platform, ECID
Ukraine: IT UKRAINE
United Kingdom: techUK

DIGITALEUROPE
Rue de la Science, 14 - 1040 Brussels [Belgium]
T. +32 (0) 2 609 53 10 F. +32 (0) 2 431 04 89
www.digitaleurope.org | info@digitaleurope.org | @DIGITALEUROPE
Transparency register member for the Commission: 64270747023-20